Stat 135, Fall 2006 A. Adhikari

HOMEWORK 4 (due Friday 9/29)

1. 8.19 **2.** 8.21a-b. **3.** 8.23. **4.** 8.27.

Exercises 5 through 10 are one big project broken into pieces. Do them in order, otherwise things won't make sense. And **for this HW, turn in your R code for each exercise.**

5. Read and thoroughly absorb Example C of Section 8.4. (Yes, that's the exercise. No, there's nothing to turn in. I'm not kidding. Read Exercise 7 and you'll see what I mean.)

6. Pick an α uniformly at random in the interval $(2, 4)$. Pick a λ uniformly at random in the interval $(1, 2)$. For the duration of this assignment, these are your own personal true values of α and λ in the gamma density. Generate an i.i.d. sample of size 200 from your gamma density, and **save the sample so that you can use it for the exercises below**. As your answer to this exercise, provide your true values of α and λ, and a histogram of your sample with the true gamma density superimposed.

7. Pretend you never knew the true values of the parameters. Compute the MOM estimates $\hat{\alpha}$ and $\hat{\lambda}$ from the sample you saved in Exercise 6.

Next, generate 1000 independent samples, each of which consists of 200 i.i.d. observations from the gamma density whose parameters are given by $\hat{\alpha}$ and $\hat{\lambda}$. Compute MOM estimates of α and λ from each sample. You now have 1000 simulated estimates of α and of λ. Draw the histogram of the distribution of each estimate, as in Figure 8.4 of your text. Comment on the shapes of the distributions (compare with Figure 8.4) and calculate the approximate standard error of each estimate. And, finally, recall that you do know the true values of α on λ. Mark each of them on the horizontal axis of the appropriate histogram. Is either of them in a surprising place?

Oh, by the way - congratulations. You are now a bootstrapper.

8. Pick two estimates of α: one at the 5th percentile of the distribution of estimates of α in Exercise 7, and the other at the 95th percentile. On a single graph, draw the histogram of your original sample and three gamma density functions: the true density, and one for each of your two chosen estimates of α (with λ in both cases equal to the MOM estimate from Exercise 7). Use three different line types to distinguish the three curves. Comment on what you see. Is it more or less what you expect, or is there something surprising?

9. Another way to bootstrap is by *resampling*. As in Exercise 7, the idea is to get the sampling distribution of your estimate by simply drawing lots of samples, computing the estimate from each sample, and noting that the observed distribution of the estimates is likely to be pretty close to the true sampling distribution. The problem lies in how to draw lots of samples when you don't know exactly what you're drawing from. Exercise 7 showed one way around this, by plugging in MOM estimates in place of the true values.

Here's another way that is simple and clever. Treat your original sample (the one you saved in Exercise 6) as a clone of the population, and draws lots of i.i.d. samples from it! That's resampling.

That is, draw 1000 independent samples, each consisting of 200 independent draws from the original sample. From each of the new samples, get MOM estimates of α and λ and draw the histograms of your estimates. As in Exercise 7, comment on the shapes, compute the approximate standard errors, and mark the true values of the parameters on the horizontal axes. Now compare with the distributions in Exercise 7.

10. Find the MLEs of α and λ. Use your MOM estimates in Exercise 7 as the initial values in the minimizing function. As your answer, provide your MLEs and the R-code you used to calculate them. Compare with the true values of α and λ.

Yes, you can bootstrap the MLEs as outlined in Example C of Section 8.5, using the same methods used for bootstrapping the MOM estimates. But you don't have to do that for this assignment.